

Towards Detecting Harmful Agendas in News Articles

Warning: This paper contains examples of online text containing strong views and offensive language.

Melanie Subbiah*
Columbia University
m.subbiah@columbia.edu

Amrita Bhattacharjee*
Arizona State University
abhattach43@asu.edu

Yilun Hua
Columbia University
yh3228@columbia.edu

Tharindu Kumarage
Arizona State University
kskumara@asu.edu

Huan Liu
Arizona State University
huanliu@asu.edu

Kathleen McKeown
Columbia University
kathy@cs.columbia.edu

Abstract

Manipulated news online is a growing problem which necessitates the use of automated systems to curtail its spread. We argue that while misinformation and disinformation detection have been studied, there has been a lack of investment in the important open challenge of detecting harmful agendas in news articles; identifying harmful agendas is critical to flag news campaigns with the greatest potential for real world harm. Moreover, due to real concerns around censorship, harmful agenda detectors must be interpretable to be effective. In this work, we propose this new task and release a dataset, NEWSAGENDAS, of annotated news articles for agenda identification. We show how interpretable systems can be effective on this task and demonstrate that they can perform comparably to black-box models.

1 Introduction

In recent years, the spread of misinformation and disinformation has become a particularly persistent and harmful issue online (Bastick, 2021; Mueller III, 2020; Vosoughi et al., 2018; Zhang and Ghorbani, 2020). For example, during the COVID-19 pandemic in the United States, we saw several instances of malicious actors propagating disinformation regarding mask mandates, vaccines, and fake remedies and cures to discredit the government and public health officials. The people initiating these disinformation campaigns typically have some harmful agenda, such as discrediting an individual/group or encouraging disruptive real-world action. Furthermore, with new conversational language models such as ChatGPT and GPT-4 (OpenAI, 2023), a malicious actor can generate human-like harmful text content at scale.

Identifying these types of harmful news campaigns typically requires consideration of three important attributes:

*These authors contributed equally to this work.

Title: Are You Prepared For The War To End All Wars?
... Everyone in the know (global elitists) knows what is happening, just not exactly when they can pull it all together. And the media awaits this war with baited breath as they count down to the dramatic moment when they can report the incident that will compel the innocent to attack the guilty. Anyone with half a brain can see the greatly increased anti-Russian propaganda of the past few weeks. This has happened as the Russia-gate claims have fallen to pieces, as former CIA analyst Raymond McGovern, the late Robert Parry, Paul Craig Roberts, and others have documented so assiduously. All across the media spectrum, from the big name corporate stenographers like The New York Times, CNN, National Public Radio, The Washington Post to The Atlantic and Nation magazines and other leftist publications such as Mother Jones and Who What Why, the Russia and Putin bashing has become hysterical in tone ...

Table 1: Example article with annotated spans from our dataset, original article from *infiniteunknown.net*, a source with label *conspiracy* in the FakeNewsCorpus. Orange spans are annotated as *conspiracy*, yellow spans are *political bias*, and blue spans are *propaganda*.

1. **Factuality** - Does the article rely on false information?
2. **Authorial Deception** - Did the author knowingly deceive the reader?
3. **Agenda** - Why did the author deceive the reader?

Misinformation in news is any article which relies on false information and can therefore be identified by focusing on *factuality*. Disinformation is deliberately misleading information created/disseminated with an intent to deceive (Shu et al., 2020a), so can be identified by *factuality* and *authorial deception*. However, the degree of harm caused by disinformation and misinformation depends on the *agenda* (or goal) of the article.

Fallis (2015) advocates for this kind of focus on agenda as a useful marker of intentionality in disinformation detection. Defining what constitutes a harmful agenda is an inherently subjective task and requires a notion of good and bad. Researchers in different domains have tried defining and formalizing the concept of harm, such as harmful online content (Scheuerman et al., 2021), COVID-19 related tweets (Alam et al., 2021), etc. However, to the best of our knowledge, the notion of harmful agendas in journalistic news articles has not been explored yet. In this paper, we therefore propose a new task of detecting harmful agendas in news articles. Inspired by definitions of harm in other works, we specifically focus on real-world harm, meaning articles that spur core belief change or actions that significantly harm someone.

To develop an initial detector, we formulate this task as classifying an article’s agenda as *harmful* or *benign*, based solely on the article text, and we annotate a dataset, NEWSAGENDAS, to evaluate performance. We note that future work could also formulate this problem in several other ways, such as also identifying the target audience, or additionally using metadata or contextual cues such as author information, publication platform, etc.

We imagine this type of agenda detector could be used to flag potentially harmful articles for further inspection. It is therefore critical that any such detector be interpretable so that further examination could quickly reveal why an article was flagged and screen out any falsely identified articles. For sensitive application areas, there is a need to build models that are interpretable by design, rather than trying to interpret their results after the fact (Rudin, 2019). Given the subjectivity and the sensitivity of this task, we build an interpretable model that uses extractive rationale-based feature detectors to ensure faithfulness and interpretability, not only at the feature level but also at the text level.

Our primary contributions are:

1. Defining the important open challenge of detecting harmful agendas in news articles.
2. Annotating and releasing NEWSAGENDAS - 506 news articles, encompassing 882 fine-grained label annotations for this task.¹
3. Developing a harmful agenda detector which jointly prioritizes interpretability and performance.

¹All data and code is available at <https://github.com/melaniesubbiah/harmfulagendasnews>.

2 NEWSAGENDAS Dataset

In order to evaluate our model’s performance and contribute an initial benchmark for this task, we annotated news articles which we are releasing as a novel dataset, NEWSAGENDAS.

2.1 Features of Interest

To promote interpretability, we hypothesize based on consultation with journalism professors at Arizona State University that the features shown in Table 2 (e.g., hate speech, propaganda, etc.) may have a significant relationship to the overall classification of article agenda in the sociopolitical context of the United States (see Table 2 for justification).

We are therefore interested in annotating these feature labels at the article-level as well as the overall agenda classification for the article. Using these features also allows us to build on the training datasets used in fine-grained news classification to classify news into these different categories.

2.2 Articles

We use articles from the FakeNewsCorpus² along with satire and real news articles from the Yang et al. (2017) dataset and propaganda articles from the Propopy corpus (Barrón-Cedeno et al., 2019) to cover a range of articles that should contain the features and agendas we are interested in. The FakeNewsCorpus contains articles in English from a web scrape of sources which frequently post misinformation. Each source has one or more specific labels indicating the general type of content it publishes and many of these labels match our features of interest (e.g., junk science, conspiracy theories, etc.). Since these labels are assigned at the source level, they serve as weak labels at the article level. We sample 600 articles for annotation, sampling to match the distribution of weak labels in the FakeNewsCorpus (based on the articles’ primary weak labels; see Appendix E for more detail).

2.3 Annotation Method

We hired Columbia University students who study journalism, political science, or natural language processing and thus have experience interpreting news (see Appendix B for hiring details).

We presented each annotator with the title of the article and the first 1,700 characters of the article truncated to the last sentence. They were asked to assume the article contained some false claims, and

²<https://github.com/several27/FakeNewsCorpus>

Label	Definition	Notes on Connection to Article Agenda
Clickbait	An exaggeration or twisting of information to shock and grab the attention of the reader.	Can be used to promote a harmful agenda (Carnahan et al.; Chen et al.), but often just a marketing strategy which is relatively benign.
Junk Science	Untested or unproven theories presented as scientific fact.	Can be unintentional, but has a high potential for harm, particularly in the medical domain (Pandey; Poynter).
Hate Speech	Language that promotes or justifies hatred, violence, discrimination, or negative prejudice against a person or category of people.	Involves extreme language that indicates clear intent on the part of the author and has a high potential for harm, even physical violence (Haynes).
Conspiracy Theory	A belief that some covert but influential organization is responsible for a circumstance or event.	Erodes public trust in science, institutions, and government (Ahmed et al.; Oliver and Wood) which may not be intentional on the part of individual actors but is harmful.
Propaganda	Promoting or publicizing a particular political cause or perspective.	Polarizes readers and harms the democratic environment necessary for healthy political debate (Guarino et al.).
Satire	Using humor, irony, or exaggeration to critique something or to amuse.	Not typically harmful when used to reveal a social/political truth, rather than for hate (Levi et al.; Golbeck et al.).
Negative Sentiment	Evokes a negative emotional response in the reader.	Evoking negative emotionality can create a lasting reaction (Mastrine), which can be more benign like sensationalism (Ward), or more harmful like negative propaganda.
Neutral Sentiment	Generally neutral/factual tone throughout the article. Does not evoke strong emotion.	Credible news organizations often have guidelines for objective and neutral reporting of 'hard-news' (Rogers).
Positive Sentiment	Evokes a positive emotional response in the reader.	Research suggests positive sentiment is not often used in disinformation or to instigate/polarize readers (Alonso et al.).
Political Bias	Angling information toward a particular political cause or perspective.	Biased articles may misrepresent/slant facts to support (harmful) agendas in cases of contentious topics (Chen et al.).
Call to Action	Urging the reader to do (or not do) something in order to further some goal.	Instigating or urging the reader to take some action for example via bandwagoning (Da San Martino et al.) may result in a (harmful) real-world effect.

Table 2: The definitions for the full set of labels annotators were asked to label articles with.

then rate whether it advanced a harmful agenda on a scale of 1 to 5. We allowed for some subjective interpretation of what a *harmful agenda* meant, but we prompted them to think of the scale of impact and whether an article might promote a real-world negative action or a strong negative belief about an individual or group of people. Lastly, they were asked to label the features found in Table 2, with the associated definitions provided, and provide 1-3 supporting evidence spans from the article for each label. They were prompted to first consider the article’s primary weak label, and not to exhaustively label features. Since the features and score were labeled separately, we did not enforce any particular relationship between an individual feature and the overall label. See Appendix C for the full task instructions. We asked them to annotate a broader list of features than we used in our models for this paper to enable future work on this problem.

The full evaluation dataset, NEWSAGENDAS, consists of 506 annotated articles with 882 fine-grained label annotations. Each article additionally has its original weak label. See Appendix D for the

label and score distribution and dataset examples.

2.4 Annotation Quality

To measure agreement between annotators, we held out an additional 90 articles for annotation by at least 2 graduate students (on average 3.4 students per article) studying natural language processing or journalism. We asked annotators just to label the harmful agenda score and to identify whether a specific feature from Table 2 was present. For each feature, we presented 5 articles with that weak label and 5 random articles. For sentiment, we presented this task as a 3-way classification between positive, neutral, and negative (see Appendix C for full task instructions). We then computed Cronbach’s alpha (a measure of internal consistency (Cronbach, 1951)) across the annotators’ responses. We observed good agreement across the harmful agenda scores (Table 3), and moderate agreement across the individual feature labels. These results indicate the data is of reasonable quality but future work could place more emphasis on how to well-annotate some of the trickier features.

Annotation Type	Cronbach’s Alpha
Harmful Agenda Scores	0.78 (0.69, 0.84)
Feature Labels	0.53 (0.35, 0.67)

Table 3: Cronbach’s alpha consistency measure for the annotated scores and feature labels in the annotation quality experiments. 95% confidence intervals are shown in parentheses. As a reference, randomly generated scores/labels produce a Cronbach’s Alpha <0.06.

2.5 Labels

We define different sets of feature labels used in the paper for clarity:

1. **Annotated gold labels** - Feature labels assigned by our annotators in NEWSAGENDAS.
2. **Weak labels** - Feature labels assigned at the source-level from the FakeNewsCorpus.
3. **BERT/FRESH labels** - Feature labels predicted by our trained models (seen in Sec. 4).

The annotated gold labels are the standard which we can evaluate our system against, but we cannot train on them since there is not enough data per label and we cannot contaminate evaluation results by training on the evaluation data. We therefore use the weak labels for training, since there is a large quantity of weak labelled articles, although they are not as accurate.

3 Methods

We leverage large weakly labeled datasets to train feature classifiers for our features of interest. We prioritize exploring different levels of interpretability in the models we compare and what performance tradeoffs come at each level. To focus our analysis, we select **7 features** to study in-depth: *clickbait, junk science, hate speech, conspiracy theories, propaganda, satire, and negative sentiment*. Out of the 4 features we excluded, 3 did not have enough labelled data. For the 4th, political bias, after consulting our journalism experts, we determined the relationship between harmful agendas in news articles and political bias is nuanced and needs further study. We therefore leave political bias to future work to promote simplicity and interpretability in our approach.

3.1 Models

As shown in Figure 1, our approach is to separately train individual neural feature classifiers for each

of the 7 features of interest. We then combine these features using a linear classifier to produce the final agenda classification. Our model is interpretable at the final level since the feature vector indicates the features that contribute to the final classification. It is also interpretable at the feature-level, where 6 of 7 features are derived from rationale-based models, which indicate the subset of input tokens that contribute to the feature classification.

Since we want to ensure faithfulness and interpretability, we derive our rationale model from the FRESH framework (Jain et al., 2020) (see Figure 1). We first finetune a BERT model (Extractor BERT) to predict a feature label from the full article text. For each token in the document, we derive a saliency score from the [CLS] token attention weights in the penultimate layer of this extractor. We extract as a rationale the top 20% most important tokens (with respect to saliency scores), irrespective of contiguity (each word is treated independently). Next, we finetune a second BERT model (Predictor BERT) to predict the feature label using only these extracted rationale tokens concatenated as input. This approach differs from the original FRESH paper in that we do not use a human-annotated dataset to introduce additional token-level supervision in rationale extraction. We also modify the FRESH framework to leverage positional embeddings for tokens. See Appendix A for details on training hyperparameters.

For the sentiment classification, we use the VADER classifier built into the NLTK Python library (Hutto and Gilbert, 2014; Loper and Bird, 2002). We choose VADER over more recent LLM-based sentiment analysis models, to facilitate interpretability. We compute the compound polarity score on a concatenation of the article title and contents. Articles with a compound score less than 0 are labeled as negative.

3.2 Training Data

For training data for the individual feature detectors, we use articles and weak labels from the same datasets described in Section 2.2 (however, we remove any articles used in NEWSAGENDAS). We handle negative sentiment labels at the model level (discussed in the next section).

Since the FakeNewsCorpus was collected from a broad scrape of unreliable websites, we noticed many of the texts did not fit the format of a news article. We therefore only use articles from the Fak-




Figure 1: The full system consists of 6 different rationale-based classifiers and a simple sentiment classifier. These 7 features are input to a linear classifier which outputs the final classification.

eNewsCorpus whose source overlaps with the list of sources used by NELA-GT (Gruppi et al., 2021) or Li et al. (2020)’s Covid-19 dataset in order to filter for high quality sources. While this approach is not exhaustive, it significantly improves the quality of the data since the sources are validated by multiple misinformation datasets. We also search and remove URLs and variants of the source names from the articles to avoid model memorization of source-label pairings.

For each individual feature detector’s training dataset, we sample 2,500 articles with the feature label we hope to detect (positive examples), and sample a range of negative examples based on a set of criteria (see Appendix E for details on negative examples for each feature). For each label, we adopt a weighted sampling strategy to increase the diversity of sources. We assign each article from a website w a weight $\frac{1}{c_w}$, where c_w is the total count of articles from website w . We then normalize these weights to sum to 1.

We additionally hold out 500 articles for the dev set and 500 articles for the test set. The test set articles come from a different set of websites than were used for the train and dev sets to make sure the test scores can not be inflated by any model memorization of website-specific styles.

4 Results

We investigate a series of research questions that analyze the efficacy of our overall approach, as well as individual components in our dataset and models.

4.1 How well can we predict the overall agenda score?

We experiment with predicting the NEWSAGENDAS annotated agenda score using different vari-

ants of our system. We fit the final logistic regression layer to the data using 10-fold cross-validation. The input is the 7 binary feature labels and the output is a binary classification of harmful or benign agenda - we bucket agenda scores 1-3 as benign and 4-5 as harmful (annotators gave a score of 3 when they were unsure of whether there was a harmful agenda in the text). We compare our method using the predicted features against three baselines: (1) predicting the majority class (0-benign), (2) using the weak source-level feature labels for logistic regression, and (3) finetuning a BERT model to classify the agenda (see Table 4). Baseline (2) demonstrates how this approach may be limited by the quality of the weak labels. Baseline (3) demonstrates a comparison against a fully black box model. We additionally compare against logistic regression using the annotated gold labels as an oracle. Using the annotated gold labels indicates a rough upper bound on performance for this type of feature-based approach, but could not be used in practice since it relies on a human annotating the articles. Note that the performance of the oracle implies a significant scope for improvement, and re-affirms our hypothesis that detecting harmful agendas in news articles is an especially difficult task for an automated system.

The oracle logistic regression model with the human annotated gold labels performs well, indicating our features of interest are very useful for the ultimate classification and promote interpretable classification of article agenda. The three systems we compare (with three different levels of interpretability) all perform better than both the majority baseline and logistic regression using just article weak labels. We also see that while we lose a little performance for every increase in interpretability (differences shown in table are statistically signif-

Method	Accuracy	Bal. Acc.
Oracle Logistic Reg.	76.7	75.6
Predict Majority Class	58.4	50.0
Weak Logistic Reg.	58.9	58.4
BERT-based Baseline	63.0 (0.21)	62.2 (0.36)
BERT System (Ours)	60.1 (0.02)	60.1 (0.02)
FRESH System (Ours)	59.3 (0.03)	59.3 (0.03)

Table 4: Overall performance (accuracy and balanced accuracy) on detecting harmful agendas in NEWSAGENDAS. The oracle logistic regression uses the annotated gold labels. Results are averaged with standard deviation shown in parentheses for neural models.

icant by a two-sample t-test, $p < .0001$), it is possible to build interpretable models that are almost as effective as the black box models for this task. This interpretability is critical because a real-world system with this accuracy would require human oversight. The strong results of the oracle model also demonstrate that investing in better feature detectors could result in improved overall agenda classification, even beyond the black-box approach.

4.2 How are the features in NEWSAGENDAS related to the overall agenda score?

We first perform a pairwise analysis of which labels are more related to higher agenda scores over others in NEWSAGENDAS, using a pairwise Wilcoxon test. *Hate speech* and *negative sentiment* are associated with higher scores most often over other labels, suggesting that these two features are particularly strong indicators of a harmful agenda. Interestingly, *call to action* loses this pairwise comparison most often, even though it seems this label would be the biggest indicator of the article encouraging a real-world outcome. This may be because *call to action* was the least represented feature in the data (only labeled 8 times) so there is not a lot of data on this feature. *Neutral sentiment* and *satire* are associated with lower scores most often over other labels, suggesting that these two features are stronger indicators of a benign agenda. See Appendix F for more details on this analysis.

We also look at the weights learned by the final logistic regression layer over the features to determine what relationship the models learn between the feature labels and the final harmful agenda score. We see that almost all of the models place the highest weight (noted in bold) on *hate speech* with *negative sentiment* and *propaganda* generally

Feature	Annot.	Weak	BERT	FRESH
Clickbait	0.96	0.12	0.47	0.45
Junk Sci.	0.22	-0.16	-0.28	-0.45
Hate Sp.	1.76	0.21	0.57	0.61
Conspir.	0.86	0.20	-0.24	-0.07
Propagan.	1.31	0.80	0.55	0.42
Satire	0.62	0.34	0.22	0.17
Negative	1.55	n/a	0.53	0.52

Table 5: Weights for each feature learned by the logistic regression models across different feature label sets. The weights are averaged across the different cross-validation subsets and across seeds when appropriate.

coming in second. The models generally place the lowest weights on *junk science*, *conspiracy theories*, and *satire*.

4.3 How well do our feature detectors work?

In order to evaluate how well each feature classifier learned its training task (predicting the weak label from the FakeNewsCorpus for its feature), we evaluate predicted labels against weak labels across three datasets: 1) the validation set, 2) the test set, and 3) NEWSAGENDAS. We compare the FRESH-based models relative to the baseline of just using the fine-tuned extractor BERT model to predict the label to explore different levels of interpretability.

In Table 6, we see that the feature classifiers generalize effectively to articles from new sources in the test set, although the performance drop (relative to the validation set) indicates that the models are relying on some source-specific qualities of articles during training. We also see reasonable performance on the articles in NEWSAGENDAS with the exception of the satire model which performs poorly. We think the poor satire performance is because the training satire articles came from higher quality websites than many of the sites in the FakeNewsCorpus and therefore the text style may be too different to transfer to many of the articles in NEWSAGENDAS.

We then evaluate how well the predicted labels agree with the annotated gold labels. To measure overlap between predicted labels and annotated gold labels, we report the intersection-over-union (IOU) and the recall for the classifiers (see Table 7). As a baseline, we include the agreement between the weak labels and the annotated gold labels. The generally low weak label agreement shows that the source-level labels for articles provide fairly distant

Feature	Val. Set		Test Set		NEWSAGENDAS	
	BERT	FRESH	BERT	FRESH	BERT	FRESH
Clickbait	90.5 (0.8)	88.7 (0.6)	61.1 (1.3)	59.0 (0.3)	76.9 (0.5)	71.6 (3.6)
Junk Science	93.9 (0.7)	93.0 (0.8)	89.3 (0.8)	89.5 (0.7)	77.4 (1.6)	73.8 (2.5)
Hate Speech	91.7 (0.3)	90.8 (0.8)	83.0 (1.3)	83.4 (0.8)	65.4 (0.7)	64.4 (1.1)
Conspiracy Theory	94.2 (0.3)	93.2 (0.6)	74.9 (1.0)	74.3 (1.6)	62.7 (1.5)	61.7 (1.4)
Propaganda	91.9 (0.6)	91.1 (0.3)	70.4 (1.1)	71.5 (2.2)	77.3 (1.3)	73.1 (1.8)
Satire	95.9 (0.2)	94.5 (0.6)	66.9 (2.1)	73.1 (2.4)	51.5 (0.4)	51.5 (0.6)

Table 6: Mean balanced accuracy scores (standard deviation in parentheses) for predicting the weak labels using the BERT and FRESH feature classifiers.

supervision relative to human judgment. We see that the BERT and FRESH models have worse but fairly similar overlap as the weak labels in many cases. The junk science and satire models have the least overlap. The black-box BERT model seems to have a slight advantage on the FRESH model, indicating there is an interpretability/performance tradeoff.

4.4 Are the extracted rationales useful?

We know that the FRESH rationales are useful to the BERT-predictors because our FRESH results show that BERT is able to achieve comparable prediction accuracy when using just the rationales as input as compared to using the entire text as input. Evaluating whether the FRESH rationales are also useful to humans is trickier. We analyze the percent of non-stopword rationale tokens that were also contained in the human-annotated rationales. However, we saw that the scores were not reliably different from just selecting the first 350 characters of the article as the rationale. This is likely because the generated rationales contain non-contiguous tokens from throughout the article, whereas the human-annotated rationales are 1-3 sentences. We therefore need to explore further human evaluation methods to quantitatively determine how well the model is rationalizing.

Through manual inspection, the rationales also seem meaningful to a human. We show three examples of common scenarios in Table 8 that demonstrate the quality of the rationales and the low word overlap score with the human-annotated rationales. The first example in this table illustrates a case where the human and FRESH model chose different labels for the article but both labels and rationales seem reasonable. The second example shows a case where the human and FRESH model agreed on the label, and the model rationale actually shares

almost all the major keywords of the human rationale (although these words are not contiguous and in the same order as in the case of the human rationale). The final example then shows a case where the human and FRESH model agreed on the label, but chose rationales with very few overlapping words other than *Washington D.C.* and *socialism*.

5 Related Work

Disinformation and Misinformation. There are many previous approaches which have studied detection of misinformation and disinformation and which would be useful in combination with the detectors developed in this work (e.g., an agenda detection system flags an article to then go through a fact-checking pipeline). Research on detecting fake news includes detectors based on linguistic features (Gravanis et al., 2019), fact-checking based systems (Ciampaglia et al., 2015), social context or propagation network based approaches (Shu et al., 2020b; Wu et al., 2015; Liu and Wu, 2018), multi-modal approaches (Khattar et al., 2019), etc. Other work has focused on characterizing/defining disinformation as a whole and developing classification schemas for campaigns (Booking et al., 2020; Fallis, 2015). However, neither disinformation detection nor characterization has explicitly looked at the more specific identification of a harmful agenda in an article.

Intent Detection. An agenda requires intention so detecting a harmful agenda is a type of intent detection. Intent detection is used in many settings with systems using slot-filling (Niu et al., 2019), conversational techniques (Larson et al., 2019; Casanueva et al., 2020), and language understanding (Qin et al., 2019). There has also been research into what intentions are involved with news articles specifically - on the intention of writing vs. sharing articles (Yaqub et al., 2020), the journalis-

Feature	IOU			Recall-1		
	Weak	BERT	FRESH	Weak	BERT	FRESH
Clickbait	32.0	30.9 (0.7)	25.3 (1.6)	53.3	46.4 (1.3)	40.2 (8.5)
Junk Science	18.5	17.1 (0.8)	12.5 (4.2)	41.7	75.0 (6.8)	77.8 (7.8)
Hate Speech	16.5	18.0 (1.3)	19.1 (2.8)	34.7	64.6 (2.5)	57.1 (7.2)
Conspiracy	27.7	18.8 (0.5)	18.4 (0.3)	40.3	60.1 (1.2)	67.5 (7.4)
Propaganda	56.2	43.0 (1.6)	40.1 (2.7)	77.1	60.0 (2.5)	59.8 (4.5)
Satire	47.9	2.8 (0.8)	2.3 (0.8)	61.4	2.9 (2.4)	2.4 (0.8)
Negative Sentiment	24.0			73.5		

Table 7: Agreement of the weak labels, BERT-predicted labels, and FRESH-predicted labels with NEWSAGENDAS’ annotated gold labels. Metric reported is mean IOU/Recall-1 (standard deviation in parentheses for predicted labels).

Human-annotated	Model-predicted	Article Opening
Negative Sentiment: American and global audiences have been bombarded with media images of wailing children in holding facilities, having been separated from adults (maybe their parents, maybe not) detained for illegal entry into the United States.	Propaganda: Atrocity Porn and Hitler Memes and Daddy ! since parents - caging children racist FDRs Indeed , voted for Trump is now Americans Nazis Hitler President Donald Trump First Lady ? sanctuary - border Trumps terrorism - His	Atrocity Porn and Hitler Memes - American and global audiences have been bombarded with media images of wailing children in holding facilities, having been separated from adults (maybe their parents, maybe not) detained for illegal entry into the United States. The images have been accompanied by gut-wrenching audio...
Junk Science: Apple cider vinegar has so many benefits, but personally one of the reasons I like it best is because of the digestive and metabolism boosting benefits.	Junk Science: 6 Things Your Body Shot Apple Cider Vinegar Daily amazing health benefits apple cider vinegar (acv salad health benefits apple cider vinegar Why ? well 1 active Metabolism It This body protein It Boosts helps iron oxygen Iron oxygen the oxygen muscle 5 It helps glucose 6 acids produce sleep a video...	6 Things That Happen To Your Body When You Take A Shot Of Apple Cider Vinegar Daily - We’ve all heard about the amazing health benefits of apple cider vinegar (acv) right? but we forget to take it or how to use it. Most people add a little bit to the top of their salad when they remember but that’s about it...
Propaganda: It isn’t any of their usual villains, the corporations and banks, who made Washington D.C. so miserable. It’s the triumph of socialism.	Propaganda: Millionaire Poverty Pimps Fight ‘Income Inequality’ Washington D.C. citys food Obama Imperial City Income Inequality in America : Oligarchy Middle Class , the town hall oligarchic city income inequality government I nt Washington Its socialism bedroom left A socialist socialist socialist nt politics millionaire poverty	Millionaire Poverty Pimps Fight ‘Income Inequality’ - The bottom fifth of Washington D.C. account for just 2% of the city’s income. It has one of the highest poverty rates in the country and the highest food stamp use. And under Obama, the Imperial City of the politicians and the poor was surrounded by some of the wealthiest districts in the country...

Table 8: Examples of labels/rationales annotated by humans and predicted by FRESH. The FRESH rationale is a concatenation of the most salient words in the text, whereas the human-annotated rationale is typically a sentence. We also highlight the FRESH-rationale words in the article opening (the title and first couple sentences) for clarity.

tic role of articles (Mellado, 2015; Tsang, 2020), and what motivates people to create and share fake news knowingly (Osmundsen et al., 2020). Finally, there has also been work on detecting deception (an intentional act) (Rubin and Conroy, 2012). However, these works have not looked specifically at automatic classification of a harmful agenda in news.

6 Conclusion

In this work, we formalize the open challenge of detecting harmful agendas in news articles, release an initial evaluation dataset, and develop an inter-

pretable system for this task. We hope our work can encourage future investment in this area - such as exploring state-of-the-art interpretable models for detecting the features we discussed, further characterizing article agenda beyond a binary classification, or investigating the interplay between text features and metadata like article source.

7 Limitations

Given the subjective nature of our proposed task, this work does have some limitations and challenges. Firstly, the notion of harm or potential

to do harm is seldom an objective factor and is also difficult to measure or quantify. Our experiments on inter-annotator agreement use a small dataset, so this study could be expanded with collaboration with social science researchers to better qualify how people perceive the agenda in different articles. Our work is also grounded in the United States, so it may have limited applications to the news in other countries (discussed more in Section 8). Secondly, our data and framework can be used to build and train a system to perform post-hoc detection of harmful agendas in news articles. However, in a real-world system, this identification would likely need to happen on the fly, so as to make readers aware of these agendas as they are exposed to the articles. Finally, another aspect that we have not addressed in this study is the effect that a platform or community may have on the perceived harm in an article. For example, on dedicated social media channels hosting discussions on alternate theories and contentious topics (such as the efficacy of COVID-19 vaccines), a junk science article with dubious claims may not be as “harmful” as opposed to the same article being posted on an open forum where readers may perceive it as scientific fact, thereby making the article more “harmful”. The context in which news articles are disseminated may have a profound impact on this perceived harm and this may be an interesting direction for future exploration.

8 Ethical Considerations

8.1 Censorship

Detecting harmful agendas in news articles has the obvious possible downstream use of filtering or banning articles which are flagged as such from being shared on social media platforms. We have already seen debate over content filtering like this take place in relation to sites like Facebook, Instagram, and Twitter moderating the dissemination of “fake news” on their platforms. One could imagine an automatic harmful agenda detector becoming part of this kind of content moderation pipeline. However, if the AI system incorrectly flags articles, it may end up censoring legitimate political speech. For this reason, we discourage any real-world use of this system at this time until further research and analysis can be completed. Additionally, we want to emphasize that this detection system should be paired with a fact-checking system to make sure that the pipeline considers the interplay between

agenda and misinformation, and does not just flag biased or opinionated free speech.

8.2 Cultural/Ideological Context

Characterizing an article as containing a harmful agenda forces definitions of what constitutes harm, which has been studied for millennia by philosophers of ethics. Normative ethics is the study of how to articulate the basic tenets of what is good and bad (Kagan, 2018). Broadly, normative ethics is divided into teleological/consequentialist (focusing on consequences to determine good/bad (Sinnott-Armstrong, 2021)) and Deontological (moral worth is intrinsic to an action (Alexander and Moore, 2021)). In this work, we focus on real-world harm which draws more on consequentialism.

Ultimately, as these opposing theories demonstrate, there is no universal interpretation of good and bad, or scale for evaluating harm. For this reason, any attempt to characterize news articles will come from a certain cultural context and perspective. The dataset we present is subject to the biases and cultural contexts of the annotators involved, so while it represents a useful starting point for work and data collection efforts in this area, future datasets around this problem must be conscious of recruiting a diverse and large annotator pool. An example of an individual bias could be that for a devout believer in the Christian God, writing which denounces God’s existence could be considered harmful disinformation. Whereas from the broader societal perspective of the United States, such a piece of writing would likely be considered a benign opinion piece.

Additionally, we want to clearly state that the framing of this research (in terms of what constitutes harm, fact, etc.) was through a United States sociopolitical context, and therefore likely does not apply across other global contexts without modifications. In conclusion, any future applications of news agenda characterization in the real world need to be very clear about the particular cultural context it is designed to operate in, what assumptions it uses, and what applications it is appropriate for.

Acknowledgments

This work is supported by the DARPA SemaFor project (HR001120C0123). The authors would also like to thank journalism subject matter experts

Kristy Roschke, Dan Gillmor and Scott Ruston for their insightful feedback and discussions. The views, opinions and/or findings expressed are those of the authors and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.

References

- Wassim Ahmed, Joseph Downing, Marc Tuters, and Peter Knight. 2020. Four experts investigate how the 5g coronavirus conspiracy theory began. *The Conversation*.
- Firoj Alam, Fahim Dalvi, Shaden Shaar, Nadir Durani, Hamdy Mubarak, Alex Nikolov, Giovanni Da San Martino, Ahmed Abdelali, Hassan Sajjad, Kareem Darwish, et al. 2021. Fighting the covid-19 infodemic in social media: a holistic perspective and a call to arms. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 15, pages 913–922.
- Larry Alexander and Michael Moore. 2021. Deontological Ethics. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*, Winter 2021 edition. Metaphysics Research Lab, Stanford University.
- Miguel A Alonso, David Vilares, Carlos Gómez-Rodríguez, and Jesús Vilares. 2021. Sentiment analysis for fake news detection. *Electronics*, 10(11):1348.
- Alberto Barrón-Cedeno, Israa Jaradat, Giovanni Da San Martino, and Preslav Nakov. 2019. Propopy: Organizing the news based on their propagandistic content. *Information Processing & Management*, 56(5):1849–1864.
- Zach Bastick. 2021. Would you notice if fake news changed your behavior? an experiment on the unconscious effects of disinformation. *Computers in human behavior*, 116:106633.
- Emerson T Booking, Alyssa Kann, and Max Rizzuto. 2020. [Dichotomies of disinformation](#). Accessed: Jan. 2022.
- Robert Carnahan, Sanjay Mishra, and Flora Teoh. 2021. [Inaccurate clickbait headline in forbes article used to promote the false claim that covid-19 vaccines change our dna](#). Accessed: Jan. 2022.
- Inigo Casanueva, Tadas Temčinas, Daniela Gerz, Matthew Henderson, and Ivan Vulić. 2020. Efficient intent detection with dual sentence encoders. *arXiv preprint arXiv:2003.04807*.
- Wei-Fan Chen, Khalid Al-Khatib, Benno Stein, and Henning Wachsmuth. 2020. Detecting media bias in news articles using gaussian bias distributions. *arXiv preprint arXiv:2010.10649*.
- Yimin Chen, Niall J Conroy, and Victoria L Rubin. 2015. Misleading online content: recognizing clickbait as "false news". In *Proceedings of the 2015 ACM on workshop on multimodal deception detection*, pages 15–19.
- Giovanni Luca Ciampaglia, Prashant Shiralkar, Luis M Rocha, Johan Bollen, Filippo Menczer, and Alessandro Flammini. 2015. Computational fact checking from knowledge networks. *PloS one*, 10(6):e0128193.
- Lee J Cronbach. 1951. Coefficient alpha and the internal structure of tests. *psychometrika*, 16(3):297–334.
- Giovanni Da San Martino, Seunghak Yu, Alberto Barrón-Cedeno, Rostislav Petrov, and Preslav Nakov. 2019. Fine-grained analysis of propaganda in news article. In *Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing (EMNLP-IJCNLP)*, pages 5636–5646.
- Don Fallis. 2015. What is disinformation? *Library trends*, 63(3):401–426.
- Jennifer Golbeck, Matthew Mauriello, Brooke Auxier, Keval H Bhanushali, Christopher Bonk, Mohamed Amine Bouzaghrane, Cody Buntain, Riya Chanduka, Paul Cheakalos, Jennine B Everett, et al. 2018. Fake news vs satire: A dataset and analysis. In *Proceedings of the 10th ACM Conference on Web Science*, pages 17–21.
- Georgios Gravanis, Athena Vakali, Konstantinos Diamantaras, and Panagiotis Karadais. 2019. Behind the cues: A benchmarking study for fake news detection. *Expert Systems with Applications*, 128:201–213.
- Maurício Gruppi, Benjamin D Horne, and Sibel Adalı. 2021. Nela-gt-2020: A large multi-labelled news dataset for the study of misinformation in news articles. *arXiv preprint arXiv:2102.04567*.
- Stefano Guarino, Noemi Trino, Alessandro Celestini, Alessandro Chessa, and Gianni Riotta. 2020. Characterizing networks of propaganda on twitter: a case study. *Applied Network Science*, 5(1):1–22.
- Suyin Haynes. 2020. [As coronavirus spreads, so does xenophobia and anti-asian racism](#). Accessed: Jan. 2022.
- Clayton Hutto and Eric Gilbert. 2014. Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 8.
- Sarthak Jain, Sarah Wiegrefe, Yuval Pinter, and Byron C. Wallace. 2020. [Learning to Faithfully Rationalize by Construction](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4459–4473, Online. Association for Computational Linguistics.

- Shelly Kagan. 2018. *Normative ethics*. Routledge.
- Dhruv Khattar, Jaipal Singh Goud, Manish Gupta, and Vasudeva Varma. 2019. Mvae: Multimodal variational autoencoder for fake news detection. In *The world wide web conference*, pages 2915–2921.
- Stefan Larson, Anish Mahendran, Joseph J Peper, Christopher Clarke, Andrew Lee, Parker Hill, Jonathan K Kummerfeld, Kevin Leach, Michael A Laurenzano, Lingjia Tang, et al. 2019. An evaluation dataset for intent classification and out-of-scope prediction. *arXiv preprint arXiv:1909.02027*.
- Or Levi, Pedram Hosseini, Mona Diab, and David A Broniatowski. 2019. Identifying nuances in fake news vs. satire: Using semantic and linguistic cues. *arXiv preprint arXiv:1910.01160*.
- Yichuan Li, Bohan Jiang, Kai Shu, and Huan Liu. 2020. Mm-covid: A multilingual and multimodal data repository for combating covid-19 disinformation. *arXiv preprint arXiv:2011.04088*.
- Yang Liu and Yi-Fang Brook Wu. 2018. Early detection of fake news on social media through propagation path classification with recurrent and convolutional networks. In *Thirty-second AAAI conference on artificial intelligence*.
- Edward Loper and Steven Bird. 2002. Nltk: The natural language toolkit. *arXiv preprint cs/0205028*.
- Julie Mastrine. 2022. [How to spot 12 types of media bias](#). Accessed: Jan. 2022.
- Claudia Mellado. 2015. Professional roles in news content: Six dimensions of journalistic role performance. *Journalism Studies*, 16(4):596–614.
- Robert S Mueller III. 2020. *Internet research agency indictment*. Good Press.
- Peiqing Niu, Zhongfu Chen, Meina Song, et al. 2019. A novel bi-directional interrelated model for joint intent detection and slot filling. *arXiv preprint arXiv:1907.00390*.
- J Eric Oliver and Thomas Wood. 2014. Medical conspiracy theories and health behaviors in the united states. *JAMA internal medicine*, 174(5):817–818.
- OpenAI. 2023. [Gpt-4 technical report](#).
- Mathias Osmundsen, Alexander Bor, Peter Bjerregaard Vahlstrup, Anja Bechmann, and Michael Bang Petersen. 2020. Partisan polarization is the primary psychological motivation behind “fake news” sharing on twitter.
- Nishant Pandey. 2021. [Fighting junk science in covid-ravaged india](#). Accessed: Jan. 2022.
- Poynter. 2021. [Several clickbait websites published misleading headlines claiming that “cigarettes block the coronavirus”](#). Accessed: Jan. 2022.
- Libo Qin, Wanxiang Che, Yangming Li, Haoyang Wen, and Ting Liu. 2019. A stack-propagation framework with token-level intent detection for spoken language understanding. *arXiv preprint arXiv:1909.02188*.
- Tony Rogers. 2019. [Objectivity and fairness in journalism](#).
- Victoria L Rubin and Niall Conroy. 2012. Discerning truth from deception: Human judgments and automation efforts. *First Monday*, 17(5).
- Cynthia Rudin. 2019. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5):206–215.
- Morgan Klaus Scheuerman, Jialun Aaron Jiang, Casey Fiesler, and Jed R Brubaker. 2021. A framework of severity for harmful content online. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW2):1–33.
- Kai Shu, Amrita Bhattacharjee, Faisal Alatawi, Tahora H Nazer, Kaize Ding, Mansooreh Karami, and Huan Liu. 2020a. Combating disinformation in a social media age. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 10(6):e1385.
- Kai Shu, Guoqing Zheng, Yichuan Li, Subhabrata Mukherjee, Ahmed Hassan Awadallah, Scott Ruston, and Huan Liu. 2020b. Leveraging multi-source weak social supervision for early detection of fake news. *arXiv preprint arXiv:2004.01732*.
- Walter Sinnott-Armstrong. 2021. Consequentialism. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*, Fall 2021 edition. Metaphysics Research Lab, Stanford University.
- Stephanie Jean Tsang. 2020. Issue stance and perceived journalistic motives explain divergent audience perceptions of fake news. *Journalism*, page 1464884920926002.
- Soroush Vosoughi, Deb Roy, and Sinan Aral. 2018. The spread of true and false news online. *Science*, 359(6380):1146–1151.
- Stephen J.A. Ward. 2010. [Emotion in reporting: use and abuse](#). Accessed: Jan. 2022.
- Ke Wu, Song Yang, and Kenny Q Zhu. 2015. False rumors detection on sina weibo by propagation structures. In *2015 IEEE 31st international conference on data engineering*, pages 651–662. IEEE.
- Fan Yang, Arjun Mukherjee, and Eduard Dragut. 2017. [Satirical News Detection and Analysis using Attention Mechanism and Linguistic Features](#). *arXiv:1709.01189 [cs]*. ArXiv: 1709.01189.
- Waheeb Yaqub, Otari Kakhidze, Morgan L Brockman, Nasir Memon, and Sameer Patil. 2020. Effects of credibility indicators on social media news sharing intent. In *Proceedings of the 2020 chi conference on human factors in computing systems*, pages 1–14.

Xichen Zhang and Ali A Ghorbani. 2020. An overview of online fake news: Characterization, detection, and discussion. *Information Processing & Management*, 57(2):102025.

A Training Hyperparameters

We use BERT-for-Sequence-Classification (bert-base-cased) from Huggingface³ for both the rationale extractor and the predictor, training on binary classification of the feature in question. We did not notice much sensitivity to hyperparameters during an initial grid-search, so we decided to use the AdamW optimizer with a learning rate of 1e-5; we applied an early stopping patience of 15 epochs and set the max number of epochs to be 50. All results are reported as an average with standard deviation across 3 different training runs (with random seeds 1000, 2000, 3000). We trained each FRESH model for several hours on 1 NVIDIA Titan Xp GPU. We also use the BERT models from the rationale extractor framework as a reference in our results since they are trained to predict the feature label from the article text. These BERT models are an artifact of training the FRESH models so they did not require additional computation.

B Annotator Recruitment and Training

We posted a recruitment notice on a journalism ListServ. We then hired the first four students who responded who met the criteria of current students at the same university as one of the authors and native English speakers. We hired the students through the university and compensated them at a rate of \$20/hour for 9-12 hours of work each. This rate is above the minimum wage in the city where the students completed the work.

After completing hiring paperwork, students had a 1-on-1 call with one of the authors who explained the goal of the research and what the task would look like, and provided a chance to discuss concerns and questions. Throughout the process students could communicate with the authors at any-time over email with questions/concerns, and they could also opt-out of the work at anytime. Otherwise students were able to complete the work independently on their own computers using the Amazon Mechanical Turk Workers Sandbox⁴. Students were compensated outside of the platform

³<https://huggingface.co/docs/transformers/index>

⁴<https://workersandbox.mturk.com>

based on their hours, and no other workers on the platform completed the tasks.

C Annotator Instructions

For the annotation of NEWSAGENDAS, students were presented with the instructions shown in Figure 2. They were not required to answer any of the questions, which allowed them to skip a whole article if the content made them uncomfortable since many of the articles contained offensive language.

Articles were displayed to the annotators as shown in Figure 3. They were then asked the questions shown in Figure 4. The feature names we used with the annotators differed slightly from the wording presented in this paper to facilitate clarity for the annotators. Whereas for this paper, we wanted to use consistent terminology throughout. Annotators could expand the label definitions in Question 2 as shown in Figure 5.

We did not ask annotators any personal or demographic questions, and neither did we collect nor store any personal information about them.

For the annotation quality experiments, students were presented with the instructions shown in Figure 6. Articles were displayed as shown in Figure 3. The students were then asked the questions shown in Figure 7 for most feature labels, but the questions shown in Figure 8 for tone-related labels. They could once again expand the definitions of the labels if needed.

D NEWSAGENDAS Label Distribution and Examples

The distribution of agenda scores labeled in NEWSAGENDAS is shown in Table 9. The distribution of weak labels, annotated gold labels, and evidence spans for each feature is shown in Table 10. We also looked at the distribution of agenda scores across each feature, which is shown as heatmaps in Figure 9. Examples of annotated evidence spans for each feature label are shown in Table 11.

E Negative Examples for Training Feature Detectors

The challenge of negative sampling arises from the potential overlaps between the class labels. For example, an article can be both "junk science" and "conspiracy theory" in practice. In the FakeNews-Corpus, the websites (and thus the articles) can have multiple labels, including a primary label that best describes the source. However, these labels

Instructions

Read the following news article carefully and then answer the series of questions that follow. Expand the full instructions for definitions and examples.

Full Instructions [\(Expand/Collapse\)](#)

You will be presented with a series of excerpts from news articles and their titles. These articles likely contain some misinformation (false, misleading, or inaccurate information). However, this task does not involve fact-checking, so please do not look anything up on the internet. Just answer the questions based off of the article as presented.

First, read each article excerpt. You will then be asked whether you feel the article advances a harmful agenda. You can provide your answer on a scale from 1-5. We understand this is a subjective question, please just answer based off of your gut reaction to the article and then spend the bulk of your time labeling the features in Question 2. In thinking about "harm", it is important to think about the scale of impact. An article which simply changes someones opinion or spreads gossip is likely less harmful than one which promotes negative actions in the real-world or strong negative beliefs about individuals or groups.

In Question 2, you will be asked more questions about specific features of the writing style and content. For each of these questions, you should copy and paste 1-3 sentences/phrases/title from the article which best support your answer into the text box. You can optionally leave the text box blank if you feel the writing style throughout the whole article supports your answer. For each feature you are asked about, you can view a definition by expanding the panel next to it. One feature will be highlighted at the top of the question - please look for this feature first.

Finally, if the article presented does not have the general format of a news article, please flag it using the checkbox and then skip it and continue on to the next article. We have tried to filter the data as well as we can but some non-article data may still be present. Only use this flag if the format looks like a conversation, tweet, etc., not just because the writing is bad or lacks professionalism. If you feel you need to skip an article for any personal reason, you can just leave the whole form blank and move onto the next article.

Please try to spend at most 5 minutes per article on average. It's okay if this means cutting back on the number of labels you are able to annotate for each article, just focus on labelling around 1-3 features per article on average.

Figure 2: The task instructions presented to annotators.

Task

Wealth of Congress: 14 vulnerable incumbents are worth at least \$1 million

WASHINGTON Fourteen vulnerable lawmakers were worth at least \$1 million at the start of this Congress. These include House incumbents and senators whose November re-election races are rated either Toss-up, Tilts or Leans by Inside Elections with Nathan L. Gonzales. Personal money isn't always advantageous in a tough campaign, but it can be helpful. Just one of these members has donated or loaned money to their campaign so far this cycle. Some of the wealthiest members who could have faced tough re-elections like Reps. Darrell Issa of California and Rodney Frelinghuysen of New Jersey aren't running. The wealthiest vulnerable lawmakers range from Missouri Sen. Claire McCaskill, whose minimum net worth ranks 24th on Roll Calls list of all current members, to Minnesota Rep. Collin C. Peterson, whose 202nd. Members need only report their financial positions in 11 broad ranges, starting with less than \$1,000 and maxing out at \$50 million or more. They do not need to report the values of their principal residences or their contents. Their liabilities during any part of 2016 are also counted. As such, our equation for calculating the minimum net worth of each member is: total minimum reported value of assets minus total minimum reported value of liabilities. 24. Sen. Claire McCaskill, D-Mo. \$26.9 million | Race rating: Toss-up Nearly all of McCaskill's assets are held by her husband, housing developer Joseph Shepard. Their only liability is \$50,000 line of credit with Enterprise Bank. The Missouri Democrat had \$9.1 million in her re-election account at the end of 2017.

Does not look like an article.

Figure 3: An example article as shown in the task format.

1. If you assume this article contains some false claims, would you say this article advances a harmful agenda?

- 1 - Definitely no 2 - Probably no 3 - I'm not sure 4 - Probably yes
 5 - Definitely yes

2. Which of the following phenomenon do you observe in the article?

In particular, pay attention to whether this article contains **propaganda**.

- Call to action ([More/Less Info](#))
 Propaganda ([More/Less Info](#))
 Neutral tone ([More/Less Info](#))
 Negative/angry tone ([More/Less Info](#))
 Positive/happy tone ([More/Less Info](#))
 Satire/humor ([More/Less Info](#))
 Clickbait ([More/Less Info](#))
 Hate/prejudiced speech ([More/Less Info](#))
 Political bias ([More/Less Info](#))
 Conspiracy theory ([More/Less Info](#))
 Junk science ([More/Less Info](#))

You must ACCEPT the HIT before you can submit the results.

Figure 4: The questions asked of annotators with an example weak label.

2. Which of the following phenomenon do you observe in the article?

In particular, pay attention to whether this article contains **propaganda**.

- Call to action** ([More/Less Info](#))

Definition: Urging the reader to do (or not do) something in order to further some goal.

- Propaganda** ([More/Less Info](#))

Definition: Promoting or publicizing a particular political cause or perspective.

Copy and paste 1-3 example sentences/phrases/title here, each on a new line and starting with 'Example:'. Or leave this area blank if you feel the whole article exhibits this feature.

- Neutral tone** ([More/Less Info](#))

Definition: Generally neutral/factual tone throughout the article. Does not evoke strong emotion.

- Negative/angry tone** ([More/Less Info](#))

Definition: Evokes a negative emotional response in the reader.

- Positive/happy tone** ([More/Less Info](#))

Definition: Evokes a positive emotional response in the reader.

- Satire/humor** ([More/Less Info](#))

Definition: Using humor, irony, or exaggeration to critique something or to amuse.

- Clickbait** ([More/Less Info](#))

Definition: An exaggeration or twisting of information to shock and grab the attention of the reader.

- Hate/prejudiced speech** ([More/Less Info](#))

Definition: Language that promotes or justifies hatred, violence, discrimination, or negative prejudice against a person or category of people.

- Political bias** ([More/Less Info](#))

Definition: Angling information toward a particular political cause or perspective.

- Conspiracy theory** ([More/Less Info](#))

Definition: A belief that some covert but influential organization is responsible for a circumstance or event.

- Junk science** ([More/Less Info](#))

Definition: Untested or unproven theories presented as scientific fact.

You must ACCEPT the HIT before you can submit the results.

Figure 5: An example of how the components in Question 2 could be expanded. Note that for 'Neutral Tone', there was no option to provide evidence as this feature was generally present throughout the article. Otherwise, if an annotator selected a checkbox, the option to provide evidence would appear.

Instructions

Read the following news article carefully and then answer the series of questions that follow. Expand the full instructions for definitions and examples.

Full Instructions [\(Expand/Collapse\)](#)

You will be presented with a series of excerpts from news articles and their titles. These articles likely contain some misinformation (false, misleading, or inaccurate information). However, this task does not involve fact-checking, so please do not look anything up on the internet. Just answer the questions based off of the article as presented.

First, read each article excerpt. You will then be asked whether you feel the article advances a harmful agenda. You can provide your answer on a scale from 1-5. We understand this is a subjective question, please just answer based off of your gut reaction to the article and then spend the bulk of your time labeling the features in Question 2. In thinking about "harm", it is important to think about the scale of impact. An article which simply changes someone's opinion or spreads gossip is likely less harmful than one which promotes negative actions in the real-world or strong negative beliefs about individuals or groups.

In Question 2, you will be asked a question about a specific attribute of the writing style and content. If you answer 'yes', you should copy and paste 1-3 sentences/phrases/title from the article which best support your answer into the text box. You can optionally leave the text box blank if you feel the writing style throughout the whole article supports your answer. For each feature you are asked about, you can view a definition by expanding the panel next to the question.

Please try to spend no more than 5 minutes per article. If you feel you need to skip an article for any personal reason, you can just leave the whole form blank and move onto the next article.

Figure 6: The annotation quality experiment instructions.

1. If you assume this article contains some false claims, would you say this article advances a harmful agenda?

1 - Definitely no 2 - Probably no 3 - I'm not sure 4 - Probably yes 5 - Definitely yes

2. Does this article contain **conspiracy theory**? [Definitions](#)

yes no

If you answered 'Yes', copy and paste 1-3 example sentences/phrases/title here, each on a new line and starting with 'Example:'. Or leave this area blank if you feel the whole article exhibits this feature.

Figure 7: The questions asked in the annotation quality experiments for most feature labels.

1. If you assume this article contains some false claims, would you say this article advances a harmful agenda?


1 - Definitely no 2 - Probably no 3 - I'm not sure 4 - Probably yes 5 - Definitely yes

2. How would you classify the tone of the article? [Definitions](#)

negative/angry neutral positive/happy

If you answered 'Yes', copy and paste 1-3 example sentences/phrases/title here, each on a new line and starting with 'Example:'. Or leave this area blank if you feel the whole article exhibits this feature.

Figure 8: The questions asked in the annotation quality experiments for tone-related labels.



(a) Counts of each harmful agenda score associated with each feature label.

(b) Fraction of the agenda scores associated with each feature label that fall into each bucket. Each row sums to 1.

Figure 9: The distribution of agenda scores associated with each feature label.

Agenda Score	Annotation Count
1	58
2	124
3	87
4	123
5	69

Table 9: Counts of the agenda scores labeled in NEWSAGENDAS. There are also 45 data points for which annotators chose not to label an agenda score but selected feature labels and evidence spans.

Feature	Weak	Annot.	Spans
Clickbait	83	110	158
Junk Science	13	15	19
Hate Speech	4	54	65
Conspiracy Theory	52	84	102
Propaganda	220	198	289
Satire	52	64	104
Negative Sentiment	–	113	103
Neutral Sentiment	–	42	–
Positive Sentiment	–	13	14
Political Bias	35	181	234
Call to Action	–	8	10

Table 10: Counts of weak labels, annotated gold labels, and evidence spans for each feature in NEWSAGENDAS.

were based on annotators’ overall impression of a website, which may not capture all possible types of its articles. Evidence suggests that websites sharing junk science articles often share conspiracy articles, or articles possessing both features (more details in the next paragraph). Then, even if a website has "junk science" as its only label, some of its articles may still be "conspiracy." Therefore, articles from this website may not be proper negative examples for a conspiracy detector.

With this observation, we develop our criteria for negative examples. For a model that detects a specific label (referred to as the positive label), we quantify the positive label’s overlap with other class labels using the overlap coefficient (Szymkiewicz–Simpson coefficient). The overlap between Label A and Label B is calculated as $\frac{|A \cap B|}{\min(|A|, |B|)}$, where A and B are the sets of websites whose multiple labels include Label A and Label B respectively. After exploratory experiments on the validation set, we adopted a threshold of 0.15 to filter out classes that overlap too much with the positive class. For example, the overlap

coefficient of "junk science" and "conspiracy" is 0.5396, exceeding 0.15. Thus, excluding "conspiracy" articles from the negative examples can better train the "junk science" detector. The negative classes after applying this criterion can be found in Table 12. In addition to sampling from these selected negative classes, all negative samples must not have the positive label among their multiple labels. Since we have multiple negative classes, we include more negative examples than positive examples, depending on the availability of the former after applying the criteria. We adopt a standard class-weighted loss in training to handle class imbalance.

F Additional Results and Analysis

The full Wilcoxon pairwise comparisons (discussed in Section 4.2) are shown in Table 10.




Figure 10: Pairwise comparisons using the Wilcoxon method across the set of features with respect to the agenda score. A *positive* Score Mean Difference with significant p-value implies that the articles with Label 1 are associated with higher agenda scores than articles with Label 2 (** $p < 0.01$, * $p < 0.05$). A *negative* Score Mean Difference with significant p-values implies the opposite. The final column indicates the Score Mean Difference. The agenda score has a bi-modal distribution, as expected in Likert scale type survey responses. Key for feature names - negprop:propaganda, callact:call to action, negemot:negative sentiment, junksci:junk science, hate:hate speech, bias:political bias, clickbait:clickbait, conspiracy:conspiracy theories, neutral:neutral sentiment, sathum:satire, posemot:positive sentiment.

Label	Example Spans from NEWSAGENDAS
Clickbait	<i>Could #RussianHackers have used a cloaking device to hide Wisconsin from Hillary?</i>
Junk Science	<i>Apple cider vinegar has so many benefits, but personally one of the reasons I like it best is because of the digestive and metabolism boosting benefits.</i>
Hate Speech	<i>They are a race of ugly dwarves, of diminutive stature, with hideous faces, evil beady eyes and stunted small minds.</i>
Conspiracy Theory	<i>The case sparked national debate over immigration reform and so-called Sanctuary Cities that shield illegals from deportation, of which San Francisco is one.</i>
Propaganda	<i>President Barack Obama made sure to shutter veterans parks in an effort to make the GOP look bad during the shutdown which occurred under his watch.</i>
Satire	<i>The former U.S. senator and former Democrat nominee for Vice President was charged with several felonies. Shockingly, felonious narcissism was not one of them.</i>
Negative Sentiment	<i>Once again, the party bereft of ideas and principle resorts to emotional obfuscation and accusation to advance their ideological prejudice.</i>
Neutral Sentiment	<i>A long lost Viking settlement known as 'Hop' is located in Canada, a prominent archaeologist has revealed.</i>
Positive Sentiment	<i>Newspapers, pamphlets and broadsheets provided nourishment to both spark the American Revolution and keep it alive.</i>
Political Bias	<i>Although this news may sound surprising, there are valid reasons for blacks to gravitate toward Trump.</i>
Call to Action	<i>We need your financial support to help reach those undecided voters, and if you would like to help, you can donate online right here.</i>

Table 11: Example evidence spans annotated in NEWSAGENDAS.

Class	Negative Example Classes
Clickbait	Conspiracy Theory, Hate Speech, Propaganda, Satire, Average
Junk Science	Hate Speech, Propaganda, Satire, Average
Hate Speech	Clickbait, Junk Science, Satire, Average
Conspiracy Theory	Clickbait, Satire, Average
Propaganda	Clickbait, Junk Science, Satire, Average
Satire	Clickbait, Conspiracy Theory, Hate Speech, Junk Science, Propaganda, Average

Table 12: All training articles belong to one of the 7 classes - Clickbait, Junk Science, Hate Speech, Conspiracy, Propaganda, Satire - or are *Average* articles, meaning likely truthful and informative. The class labels are from the FakeNewsCorpus, Propopy Corpus, and the Yang et al. (2017) satire dataset. We omit articles from websites that only have less informative labels such as *bias* or *political*.